

Kent Academic Repository

Full text document (pdf)

Citation for published version

Vinson, David and Ponari, Marta and Vigliocco, Gabriella (2013) How does emotional content affect lexical processing? In: Annual Meeting of the Cognitive Science Society, 31/7 - 3/8/2013, Berlin, Germany.

DOI

Link to record in KAR

<http://kar.kent.ac.uk/44826/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

How does emotional content affect lexical processing?

David Vinson, Marta Ponari & Gabriella Vigliocco

{d.vinson, m.ponari, g.vigliocco}@ucl.ac.uk

Cognitive, Perceptual and Brain Sciences, University College London
Gower Street, London WC1E 6BT United Kingdom

Abstract

It is now generally accepted that words' emotional content plays a role in lexical processing, but the literature offers incompatible findings concerning what this role may be. Here we use a large sample of lexical decision data (British Lexicon Project, Keuleers et al., 2012) and we carry out a series of analyses differing in the way emotional variables are treated. A variety of statistical approaches yielded common conclusions: when confounding variables are taken into account, emotional words, whether positive or negative, are processed faster than neutral words. This effect is categorical rather than graded; is not modulated by emotional arousal; and is not limited to words explicitly referring to emotions. We discuss this in terms of internally grounding words' meanings in emotional experience, akin to the manner in which concepts may be grounded in perception and action.

Keywords: emotion; valence; lexical decision.

Introduction

In mainstream lexical processing studies, emotional content has been largely ignored, whether considered irrelevant to the core meanings of words, or as properties of narrowly defined sets of words explicitly referring to emotions (e.g. Altarriba & Bauer, 2004). Recently, however, a number of studies of lexical processing effectively demonstrated that emotional content plays a role even in shallow tasks involving single words such as lexical decision (e.g. Estes & Adelman 2008a,b; Kousta, Vinson & Vigliocco, 2009; Kousta, et al., 2011; Larsen, et al., 2008).

As a result, language processing researchers have begun to acknowledge the interplay between emotion and language processing systems, discussing emotional effects in language processing as due to the embodied nature of linguistic representations (e.g. Kousta et al., 2011; Moseley, et al., 2012; Vigliocco et al., 2009), just as researchers in other domains of cognition have posited embodied emotional effects (e.g. Pistoia et al., 2010). However, precisely which mechanisms are involved in emotional processing is unclear at the present. This is because different studies of lexical processing have found different and apparently incompatible results even when the same task (e.g., lexical decision) is used.

It has been shown that previously reported effects of emotional valence (i.e. numeric ratings indicating the extent to which a word is positive, neutral or negative) can change dramatically once confounding variables such as length, frequency and orthographic neighbourhood size are taken into account (Larsen, Mercer & Balota, 2006). However, even after controlling for non-emotional variables, results are conflicting. Estes & Adelman (2008a,b) and Larsen et al (2008) reported slower lexical decision reaction times (RTs)

for negative than positive words. This has been interpreted in terms of attentional vigilance: heightened and/or extended attention to negative stimuli (e.g. Fox et al., 2001; Pratto & John, 1991) which would slow any decision (such as lexical decisions) on other aspects of the stimuli. In contrast, Kousta et al. (2009) found a processing advantage for both negative and positive over neutral words, which they explain in terms of greater motivational relevance of emotionally loaded stimuli (e.g. Lang, Bradley & Cuthbert, 1997). Kousta et al argued that the discrepancy in findings was due to the relative lack of neutral words in the data sets tested previously, but especially due to the lack of control of potentially confounding variables, such as ratings of familiarity and age of acquisition (AoA) in previous studies.

In addition, Larsen et al (2008) found that the effect of valence was modulated by the arousal of words such that a negative disadvantage was present for medium-low arousing words, but no effect was observed for highly arousing negative words, Estes and Adelman (2008b) argued for a far more constrained role of arousal, and Kousta et al. (2009) argued that valence effects could not be explained in terms of arousal (although these authors did not explicitly test for valence \times arousal interactions).

All of these previous studies were conducted using lexical decision data from a single source: the English Lexicon Project (ELP, Balota et al., 2007), so in addition to questions about the different assumptions and approaches taken by previous authors, one may also wonder about the extent to which the findings may be related to quirks of that particular item set. Here we take advantage of an entirely independently obtained large-scale set of lexical decision data (British Lexicon Project (BLP): Keuleers, Lacey, Rastle & Brysbaert, 2012), to try and resolve these questions. Our analyses compare models based on different *a priori* theoretical assumptions concerning the role of valence in word processing, controlling non-emotional variables known to affect lexical decision RTs. We begin by fitting baseline models in which all the non-emotional predictors mentioned above are taken into account, then add specific terms embedding different assumptions about the role of valence. Such an approach is essential in order to test theoretical accounts of emotion effects in lexical processing.

After assessing how well different measures of valence perform after taking baseline variables into account, we move on to evaluating the role of other aspects of emotional content besides just valence, assessing the extent to which valence effects may instead be explained or modulated in terms of arousal. Finally, we test whether the effects of emotional valence differ for words specifically referring to emotional experience, vs. words that are only valenced.

Materials and methods

Data

From the full set of words in the BLP, we selected those 1374 words for which valence ratings were available from ANEW (Bradley & Lang, 1999), or from the additional ratings described in Kousta et al., 2009, 2011). Next, we filtered out those words for which BLP participants were extremely inaccurate on making lexical decisions: those with overall accuracy less than 67% in the BLP ($n=56$, e.g. *larkspur*, *dryad*, *godhead*). This is an important step not employed in previous studies, as widely unfamiliar words are likely to elicit slow RTs, and to receive neutral valence ratings from participants. Finally, we removed five words for which concreteness and imageability ratings were not available, leaving 1313 words for analysis. Of these, 856 were in common with the set from the ELP which Kousta et al. (2009) analysed.

Measures of emotional valence

We centred the scale of the original valence ratings which ranged from 1-9, so as to range from -4 (most negative) to +4 (most positive) with 0 reflecting neutrality. We then created the following measures that embed different theoretical assumptions concerning valence. The most essential distinction concerns the direction of valence effects as this differentiates between highly distinct accounts of emotion processing. If the crucial distinction is between negative words and other words, this would favour attentional vigilance or other negativity-bias accounts of emotion processing; but if instead the crucial distinction is between emotionally valenced and neutral words, it would favour motivational accounts of emotion processing.

In addition, we compare models in which valence is considered as a continuous measure, vs. models in which it is discretized, as a test of previous claims that effects of emotion should be considered all-or-nothing (e.g. Estes & Adelman, 2008a,b).

Continuous valence These measures treat valence as a graded measure varying from most negative (-4) through neutral (0) to most positive (+4).

Linear measure includes only the linear relationship between valence and RT. If negative words are slower than other words (e.g. Estes & Adelman, 2008a,b; Larsen et al., 2008), we expect to find a negative slope (RTs decrease with increasing valence).

Polynomial measure includes linear and quadratic components of valence.¹ If valenced words are faster than neutral words with no difference between positive and

negative (e.g. Kousta et al., 2009) we expect a negative quadratic coefficient while the linear coefficient would offer no further benefit.

Discrete valence These measures treat valence as categorical rather than continuous/graded.

Negative/positive measure includes two discrete valence classes: negative (valence <0) and positive (valence ≥ 0) valence levels. If negative words are slower than other words, these two categories should differ. This model is the simplest discrete counterpart to the linear measure above, and was preferred by Estes and Adelman (2008a) as more complex measures tested did not account for the data better.

Valenced/neutral measure treats positive and negative as a single class, compared to neutral (emotional: $|valence| > 1.5$; neutral: $|valence| \leq -1.5$). If emotional words are faster than neutral words we expect to find differences between these two categories (as we would for the quadratic term of the polynomial measure).

Design and analysis

We fit a variety of linear mixed-effects models described in more detail below, in each case testing for a partial effect of valence on lexical decision latencies, using any of the four proposed valence measures. We conducted our analyses on log-transformed RT (excluding error trials) then replicated the same analyses on untransformed RT to be sure that log-transformation did not produce anomalous patterns of results. Additionally, we fit models not only to trial-level data but also to item averages, for both log(RT) and untransformed RT.²

Analysis of trial level data was carried out using linear mixed-effects models (packages lme4: Bates & Maechler, 2009; and languageR: Baayen, 2009; cf. Baayen, 2008) in the R programming environment (R Development Core Team, 2009). Model fits included random intercepts for both subjects and items, as well as random slopes by subjects (for emotional predictors only, which are constant for each item). Analysis of item averages was carried out using ordinary least squares regression. Below we focus upon the results for trial-level analyses of log(RT) but across the board the findings are comparable for analyses of untransformed RT and/or item averages.

In all of the analyses we conduct upon valence measures, we always begin with a baseline model in which we consider the following non-emotional factors that were

¹ Kousta et al (2009) modelled nonlinearity using restricted cubic splines; here we report a measure including linear and quadratic terms because they map directly onto the theoretical alternatives described in previous literature. We also tested models based on restricted cubic splines; they perform comparably to the polynomial models described above.

² Some previous studies of this nature only report analyses on average response times for single words, averaged across multiple subjects but treated as point estimates (e.g. Estes & Adelman, 2008a,b; Kousta et al., 2009; Larsen et al., 2008). Such approaches may overestimate the quality of any predictor as an essential component of variability has been discarded. In the present study we conduct analyses of trial-level data (nearly 50,000 observations) as well as item averages, allowing us to test whether emotional variables still play a role when individual variability is taken into account.

controlled in all the previous studies we have mentioned: number of letters; log(HAL frequency), orthographic neighbourhood size (all from Balota et al., 2007); We also included additional non-emotional predictors controlled by Kousta et al. (2009) and which those authors argued to be essential in order to unambiguously interpret effects as emotional in nature: mean positional bigram frequency (Balota et al., 2007); ratings of concreteness, imageability and familiarity (Coltheart, 1981) and age of acquisition ratings (Stadthagen-Gonzalez & Davis, 2006). As a result, our tests of emotional variables provide results that can be unambiguously attributed to emotion rather than other characteristics of words with which emotional properties may be confounded.

The role of arousal

Some previous studies have shown that effects of valence are modulated by arousal (Estes & Adelman, 2008b; Larsen et al., 2008, but see Kousta et al. 2009). Using a similar modelling approach as above, we test the role of arousal in two ways. First, we treat arousal as a categorical measure (high arousal words vs. low arousal words), testing valence \times arousal interactions for any of the valence measures described previously which turn out to be significant predictors of lexical decision RT. If arousal modulates the effect of valence we should see such an interaction. Second, we treat arousal as a control variable, testing in a different set of models whether unique effects of valence can be observed after variation related to arousal is taken into account. This is particularly important for models distinguishing valenced from neutral words (i.e. quadratic term of the polynomial measure, and valenced/neutral measure) as valenced words also exhibit a strong tendency to be more arousing as well (Bradley & Lang, 2000).

Emotion words vs. emotionally valenced words

One essential factor that has been neglected so far in large-scale studies of emotion in lexical processing is whether any valence effect is being driven by a specific, limited set of words: those referring explicitly to emotion (e.g. *love, shame, joy, hate* in contrast to valenced words not directly referring to emotions, e.g. *prison, cake, justice, cheat*). For example, Altarriba and Bauer (2004) argue that emotion words are sufficiently different to other types of words that we ought to consider words as falling into three categories: concrete, abstract and emotion words. Moreover, it has been argued that emotion words may be embodied not only internally (via emotional experience, as argued by Kousta et al., 2011) but also due to body states associated with emotional experience (such as facial expression, posture etc., Moseley, et al., 2012). Such words tend to be prevalent in our vocabulary and even a cursory inspection of valence norms reveal many such items among the set. If these words alone are responsible for emotion effects, one cannot conclude that valence is relevant to lexical processing in general, as it may play a role only in the specific, tightly constrained domain of emotion words.

To address this issue, we used Wordnet-Affect (Strapparava & Valitutti, 2004) to identify emotion words. Wordnet-Affect classifies words according to their organisation in Wordnet. Any word with an emotional sense is considered "emotional", thus this is a conservative classification. We hand-classified a few additional words as potentially emotional (e.g. *courage, craven, stern*) with 193 of the 1313 words classified as emotion words. To test whether emotion words alone are responsible for valence effects we fit models as above, testing the interaction between valence and emotion-word classification. If emotion words drive the effects observed we should see an interaction such that the valence effects are restricted to emotion words (or at least differ between emotion and non-emotion words).

Results

Fitting baseline models

It is no surprise that many of the non-emotional variables were significant predictors of lexical decision latencies, consistent with a wealth of previous studies. For the purposes of the present study we simply note here that higher-order polynomial transformations offered significant improvement in performance over linear-alone components for several of the predictors. Moreover, although some factors were not significant predictors in the baseline model (i.e., concreteness, imageability and summed positional bigram frequency) we retained them as (linear) predictors³ along with the following predictors that were significant in the (reduced) baseline model: 3-order polynomial transformations: (log frequency, number of letters, number of orthographic neighbours, familiarity); linear terms (age of acquisition) (see Figure 1).

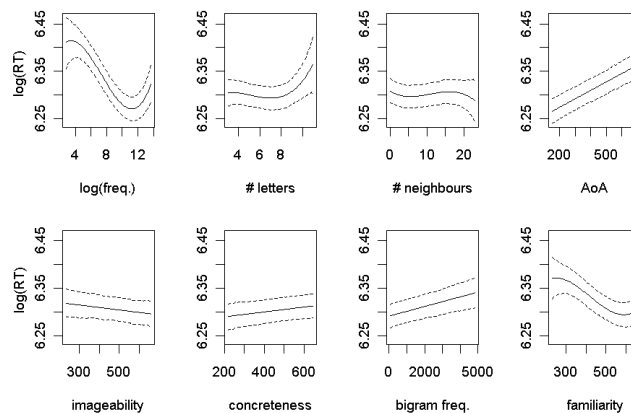


Figure 1: Predictors in the baseline model (logRT; trial-level data). Dashed lines depict 95% highest posterior density CI (parameter estimates). Similar patterns were observed for item-level analyses and untransformed RT.

³ Non-significant predictors were kept in the baseline model in case their absence may have altered the effects of emotional valence in subsequent models.

Measures of emotional valence

We tested the effects of valence by adding each of the valence measures described above to the best-fit baseline model above, thus always allowing us to evaluate the partial effect of valence only after non-emotional variables were taken into account. We also added that same valence term in each model as a random slope by subjects (in analysis of trial-level data).

Table 1: Partial effects of the different valence measures (logRT, trial-level), taking non-emotional variables into account. The same patterns were observed for item-level analyses and for untransformed RT

Valence measure	Estimate (Std err.)	<i>t</i> statistic
Linear	.00066 (.00161)	0.57
Polynomial ⁴		
(linear term)	.00052 (.00116)	0.45
(quadratic)	-.00158 (.00076)	-2.07
Negative/positive	.00166 (.00286)	0.58
Valenced/neutral	-0.0067 (.00357)	-1.87

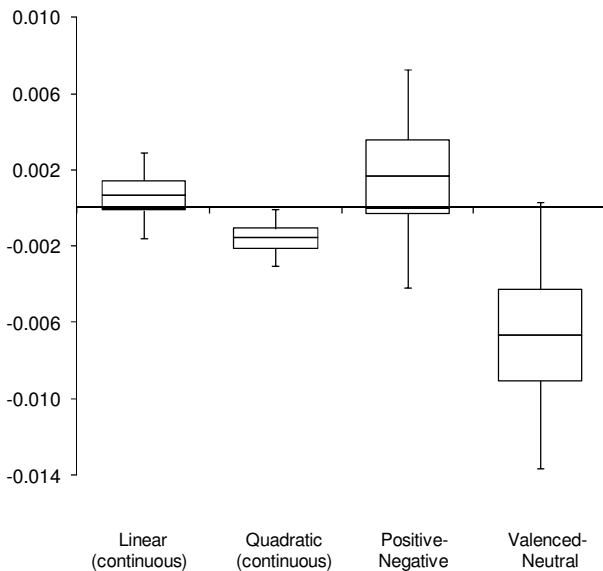


Figure 2: Graphical depiction of parameter estimates of the different valence measures reported in Table 1. In all cases a value of 0 indicates no effect. For continuous measures, the values plotted reflect the slope (linear measure) or quadratic coefficient (polynomial measure). For categorical measures they represent the estimate of the difference between the two conditions (logRT). Horizontal line = mean parameter estimate. Box depicts 50% confidence interval of the parameter estimate; whiskers depict 95% confidence interval.

Only those measures in which valenced words differ from

⁴ We also tested 3-order polynomial measure of valence, but the cubic term was never a significant predictor in any of our analyses.

neutral words (quadratic term of polynomial continuous measure, and valenced/neutral measure) were reliable predictors of lexical decision RT. Instead, the linear continuous measure did not predict RT once confounding factors were taken into account, and the same was true for the negative/positive measure. To assess whether the continuous (quadratic) measure offers sufficient additional explanatory power beyond the simplest categorical measure contrasting valenced to neutral words, we fit one additional set of models, in which we entered 2-order polynomial valence along with the categorical measure, and we compared the models using likelihood ratio tests. There was no significant improvement gained by adding this additional term (log likelihood ratio for valenced/neutral model = 7629.1; log likelihood ratio for combined model = 7630.0; $\chi^2(5) = 1.7972$, $p = .876$ with comparable results for item-level analyses and analyses of untransformed RTs).

At this stage the data suggest that the effect of valence is best described as a simple, categorical contrast between words with emotional associations and those without. Thus, when non-emotional variables are taken into account, we see that a categorical measure of valence, regardless of polarity, is sufficient to account for emotional effects in word processing.

The role of arousal

Here we focus upon those valence measures which were reliable predictors in the previous section (i.e., 2-order Polynomial and Valenced/Neutral), assessing whether they can be accounted for, or modulated, by arousal.

First, we tested the interaction between arousal and each of the two valence measures (continuous and categorical). For these analyses we discretized arousal, using a median split to characterise words as low or high arousal (contrast coded). For trial-level analyses we included both main effects and the interaction as random slopes by subjects. We found that the main effect of valence persisted, with no effect of arousal category and no interaction between the two. For log RT and trial level analysis: quadratic coefficient estimate = $-.00338$ (SE = $.00110$), $t = -3.067$, arousal main effect and interaction $|t| < 1.2$; categorical coefficient estimate = $-.0135$ (SE = $.0050$), $t = -2.725$, arousal main effect and interaction $|t| < 1$ (with item level analyses and analyses of untransformed RTs showing the same pattern).

Next, we instead considered arousal as a continuous measure of arousal into the models, testing whether a partial effect of a valence measure could still be seen after arousal was taken into account. For trial-level analysis this meant including random slopes by subject for arousal as well. We started by adding arousal to the baseline model described above. When arousal was the only emotional variable included, its effects were significant (estimate of the slope = $-.0050$ (SE = $-.0020$, $t = -2.518$): more arousing words elicited faster responses. We then added a valence measure to this baseline+arousal model. For both the polynomial and the categorical valence measure, effects persisted once

arousal was taken into account. For log RT and trial level analysis: quadratic coefficient estimate = -.00261 (SE = .00144), $t = -2.627$; categorical coefficient estimate = -.0112 (SE = .0046), $t = -2.410$, with the partial effect of arousal not reaching significance in either case ($|t| < 1$), a finding replicated in item-level analyses and analyses of untransformed RTs. These effects of emotion can thus be attributed to valence rather than arousal.

Emotion words vs. emotionally valenced words

As in the second set of analyses considering the role of arousal, we tested whether the effects of valence described above were different for emotion words and those not referring to emotional states (using Wordnet-Affect, Strapparava & Valitutti, 2004), by testing for statistical interactions.

Just like our analyses involving arousal, the main effect of valence was unchanged, with no effect of Wordnet-Affect category and no interaction. For log RT and trial level analysis: quadratic coefficient estimate = -.00197 (SE = .00085), $t = -2.31$, Wordnet-Affect category main effect and interaction $|t| < 1$; categorical coefficient estimate = -.00969 (SE = .00419), $t = -2.31$, Wordnet-Affect category main effect and interaction $|t| < 1.02$ (again with item level analyses and analyses of untransformed RTs showing the same pattern).

Discussion

Our analyses show a reliable, consistent and rather simple pattern of emotion effects in lexical processing. Once potentially confounding variables are taken into account, lexical decisions to emotionally valenced words are recognised faster than those to neutral words. This finding differs from some previous studies (Estes & Adelman, 2008a,b; Larsen et al., 2008): those investigating ELP data, using a more limited set of words (from ANEW, Bradley & Lang, 1999) and crucially, for which some important control variables are unavailable. Those studies also conducted analysis over item averages only, allowing the possibility that valence effects observed there may have been magnified or distorted as a consequence of treating these values as point estimates rather than varying by subjects. Our results also differ from those reported by Kousta et al. (2009) although consistent with their overall conclusions. There appears to be no benefit in considering valence as a continuous measure: the 2-order polynomial valence model is no better than the simplest categorical model (valenced vs neutral). In fact when we reanalyze their data set, there too we find that a categorical measure contrasting valenced to neutral words is comparable to the continuous measure they favoured.

We also found this categorical effect of valence to be general in nature: it is not modulated by arousal, and it is not driven by words specifically referring to emotional experience. This finding resonates with recent neuroimaging evidence using a highly controlled set of words, in which activation in rostral anterior cingulate cortex (an area

associated with emotion processing) is modulated by valence (regardless of whether it is positive or negative) and not by arousal (Vigliocco et al., 2013).

Why would emotional content, regardless of polarity, facilitate lexical processing? Under general motivational accounts of processing (Lang, Bradley & Cuthbert, 1997) both negatively and positively valenced items are relevant to survival and wellbeing albeit for different reasons. Crucial in this regard is the involvement of emotion processing systems even for lexical stimuli which do not exhibit obvious low-level visual characteristics argued to be evolutionarily linked to positive or negative emotions (vs. emotional expressions or visual properties of dangerous entities). In a recent proposal, the involvement of emotional systems has been argued to provide a means for grounding abstract concepts in internal experience, just like concrete concepts are accepted to be grounded in sensory-motor experience (Kousta, Vigliocco, Vinson, Andrews & Del Campo, 2011; Vigliocco, Meteyard, Andrews & Kousta, 2009). Under this view, emotional content of words would facilitate their processing, in a manner akin to the way in which sensory experience (operationalised as imageability or concreteness; Kousta et al. 2011) facilitates processing.

Acknowledgments

Supported by UK Economic and Social Research Council grant RES-062-23-2012 to Gabriella Vigliocco.

References

- Altarriba, J., & Bauer, L. M. (2004). The distinctiveness of emotion concepts: A comparison between emotion, abstract, and concrete words. *American Journal of Psychology*, 117, 389-410.
- Baayen, R. H. (2008). *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge University Press.
- Baayen, R.H. (2009). languageR: Datasets and functions with "Analyzing Linguistic Data: A practical introduction to statistics". <http://cran.r-project.org/package=languageR>.
- Baayen, R.H., Davidson, D.J. and Bates, D.M. (2008) Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390-412.
- Balota, D.A., Yap, M.J., Cortese, M.J., Hutchison, K.A., Kessler, B., Loftis, B., Neely, J.H., Nelson, D.L., Simpson, G.B., & Treiman, R. (2007). The English Lexicon Project. *Behavior Research Methods*, 39, 445-459.
- Bates, D. & Maechler, M. (2009). lme4: Linear mixed-effects models using S4 classes. <http://cran.r-project.org/package=lme4>.
- Bradley, M.M. & Lang, P.J. (1999). *Affective Norms for English Words (ANEW): Stimuli, instruction manual and affective ratings* (Tech Report C-1). Gainesville: University of Florida, Center for Research in Psychophysiology.

- Bradley, M. M., & Lang, P. J. (2000). Measuring emotion: Behavior, feeling, and physiology. *Cognitive Neuroscience of Emotion*, 25, 49-59.
- Coltheart, M. (1981). The MRC Psycholinguistic Database. *Quarterly Journal of Experimental Psychology*, 33A, 497-505.
- Estes, Z., & Adelman, J. S. (2008a). Automatic vigilance for negative words in lexical decision and naming: Comment on Larsen, Mercer & Balota (2006). *Emotion*, 8, 441-444
- Estes, Z., & Adelman, J. S. (2008b). Automatic vigilance for negative words is categorical and general. *Emotion*, 8, 453-457.
- Fox, E., Russo, R., Bowles, R. & Dutton, K., (2001). Do threatening stimuli draw or hold visual attention in subclinical anxiety? *Journal of Experimental Psychology: General*, 130, 681-700.
- Keuleers, E., Lacey, P., Rastle, K., & Brysbaert, M. (2012). The British Lexicon Project: Lexical decision data for 28,730 monosyllabic and disyllabic English words. *Behavior Research Methods*, 44, 287-304
- Kousta, S.-T., Vigliocco, G., Vinson, D.P., Andrews, M. & Del Campo, E. (2011). The representation of abstract words: Why emotion matters. *Journal of Experimental Psychology: General*, 140, 14-34.
- Kousta, S.-T., Vinson, D. & Vigliocco, G. (2009). Emotion words, regardless of polarity, have a processing advantage over neutral words. *Cognition*, 112, 473-481.
- Lang, P.J., Bradley, M.M. & Cuthbert, B.N. (1997). Motivated attention: Affect, activation and action. In P.J. Lang, R.F. Simons & M.T. Balaban (Eds.), *Attention and orienting: Sensory and motivational processes*. Hillsdale NJ: Erlbaum.
- Larsen, R.J., Mercer, K.A. & Balota, D.A. (2006). Lexical characteristics of words used in emotional Stroop experiments. *Emotion*, 6, 62-72.
- Larsen, R.J., Mercer, K.A., Balota, D.A. & Strube, M.J. (2008). Not all negative words slow down lexical decision and naming speed: Importance of word arousal. *Emotion*, 8, 445-452.
- Moseley, R., Carota, F., Hauk, O., Mohr, B. & Pulvermüller, F. (2012). A role for the motor system in binding abstract emotional meaning. *Cerebral Cortex*, 22, 1634-1647.
- Pistoia, F., Conson, M., Trojano, L., Grossi, D., Ponari, M., Colonnese, C., Pistoia, M.L., Carducci, F & Sarà, M. (2010). Impaired conscious recognition of negative facial expressions in patients with locked-in syndroms. *Journal of Neuroscience*, 30, 7838-7844.
- Pratto, F. & John, O.P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology*, 61, 380-391.
- Stadthagen-Gonzalez, H., & Davis, C.J. (2006). The Bristol Norms for Age of Acquisition, Imageability and Familiarity. *Behavior Research Methods*, 68, 598-605.
- Strapparava, C. & Valitutti, A. (2004). WordNet-Affect: An affective extension of WordNet. *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004)*, 1083-1086.
- Vigliocco, G., Meteyard, L., Andrews, M., & Kousta, S.-T. (2009). Toward a theory of semantic representation. *Language and Cognition*, 1, 219-247.
- Vigliocco, G. Vinson, D., Della Rosa, P., Kousta, S.T., Devlin, J. & Cappa, S.F. (2013) The neural representation of abstract words: A role of emotion. In press, *Cerebral Cortex*.